

SKIF-AURORA



PSI RAS



RSC SKIF



SUSU



“SKIF-Aurora” solution has been developed under the framework of “SKIF-GRID” supercomputing project of the Union of Russia and Belarus by Ailamazian Program Systems Institute of the Russian Academy of Sciences, RSC SKIF company and other “SKIF-GRID” consortium collaborators in alliance with Eurotech and with technical support from Intel Corporation

SKIF-AURORA



PSI RAS



RSC SKIF



SUSU



*I am easily satisfied with the very best
Winston Churchill*

USE SKIF-AURORA. EASY. THE VERY BEST.



“SKIF-AURORA” OVERVIEW

“SKIF-Aurora” is a scalable supercomputer system that has been developed under the frame-work of supercomputer project “SKIF-GRID” of the Union of Russia and Belorussia. The “SKIF-GRID” is a government initiative of Russia and Belorussia Union, targeted to boost technology development in high-performance computing and grid technologies (<http://skif-grid.botik.ru/>).

“SKIF-Aurora” key objective is to deliver peta- and exa-flops-scale solutions for supercomputers for the world most challenging scientific problems, while utilizing mostly commodity building blocks combined with own critical know-how.

The “SKIF-Aurora” project is conducted by broad consortium of collaborators in Russia and abroad. Principal collaborators include:

- RSC SKIF — HPC system integrator and solution provider for the new level super-computing;
- Ailamazyan Program Systems Institute of the Russian Academy of Sciences (PSI RAS) — “SKIF GRID” project lead organization and co-developer of key technologies and subsystems for “SKIF-Aurora”.

The first “SKIF-Aurora” installation has been landed in 2009 in South Ural State University (SUSU), which is aarded the status of Russian National Research University.



ABOUT “SKIF-GRID” PROJECT



Supercomputer project “SKIF-GRID” of the Union of Russia and Belorussia has been initiated in April 2007 by government decree, with 2007–2010 execution time-frame. The project goal is to foster and support of original technology development in high-performance and grid computing field. The project includes four major directions: grid technology, high-performance computing, information security and pilot applications. Project results include development of “SKIF” series 3 and series 4 supercomputing families and establishment of “SKIF-Polygon” federation of HPC centers and grid system. Lead organizations from the Russian and Belarus sides are respectively Ailamazyan Program Systems Institute of the Russian Academy of Sciences (PSI RAS) and United Insitite of Informatics Problems of National Academy of Sciences of Belorussia (UIIP NASB). More information is available in web sites

<http://skif-grid.botik.ru/> and <http://skif.bas-net.by>



“SKIF-AURORA@SUSU” PARAMETERS

- Number of nodes/CPUs/cores: 256/512/2048
- CPU type: Intel Xeon X5570 (four-core, 2.93 GHz)
- RAM: 3 TB (DDR3)
- Storage: 60 TB of external storage, 20,5 TB SSD on nodes
- Interconnect: 3-D torus topology (60 Gbit/s per node max., latency 1 μ s)
- Auxillary network: InfiniBand QDR (40 Gbit/s, latency 2 μ s)
- Global synchronization network
- Three redundant monitoring and management networks, including SKIF-ServNet sensor and management network
- Programmable FPGA accelerators
- Peak (theoretical) performance: 24 TFlops
- LINPACK performance: 21.84 Tflops

Energy efficiency

- Effective cooling: chip-level liquid cooling enables more efficient heat transfer in comparison with forced air cooling
- TCO: up to 60% saving in cooling electricity cost

Compatibility

- Fully compatible with x86 architecture
- Optimized MPI-2 libraries

Reliability

- No moving components in computational part
- Three-layer monitoring and management solution
- All major systems redundancy
- Hot-swappable system components.



“SKIF-AURORA” SOLUTION BENEFITS

We present solution for unique state-of-art energy-efficient data-centers to Pflops scale and beyond:

- Best effectiveness in terms of energy consumption and cooling—green design
 - Up to 60% savings in electricity consumption in the installation and over the years
 - Free cooling throughout the year is an important part of cost-saving
- Highly possible compute density for non-accelerator-based x86 solution
 - The highest density of computational power for standard x86-64 processor system – more than 10 CPU each up to 130W TDP per 1U
 - With six-core Intel Xeon X5600 processors — up to 40 Tflops peak per rack has been delivered
- Unique High bandwidth interconnect — up to 60GB/s
- Improved system scalability
 - Achieve 1 Pflops in just 25 racks of “SKIF-Aurora” consuming about 2,84 MWatt
- Unique 3-D torus interconnect topology enables highest scalability of the system to multi-petaflops
- Dedicated network for improved implementation of synchronization of clustered jobs, and reduction of OS jitter noise
- Improved effectiveness of collective operations — interconnect NIC and CPU offload engine
- Effective support of non-MPI programming models for HPC — PGAS, GASNET, etc.
- Full compatibility with x86 platform software while ability to utilize FPGA for offload and acceleration
- Supercomputer highest reliability
- Effective power supply system
- Improved ergonomic and aesthetic features
- Unique monitoring and management system with full-size LCD touch screens on every chassis
- Support of large research consortium in Russia and Belarus with involvement in R&D alliance with Eurotech S.p.A.
- Full compatibility with HPC software stack of commodity x86 cluster systems, combined with a broad collaborative activity in software porting, tuning and optimization by “SKIF-GRID” consortium.



KEY “KNOW-HOW” OF RUSSIAN PARTNERS

Unique power supply and cooling system with state-of-the-art PUE (Power Usage Effectiveness). DC power system distribution provides opportunities in:

- More efficient AC/DC power conversion (to 48V)
- Savings in uninterruptible power supply design due to removed conversions from out-battery DC back to AC

Liquid cooling infrastructure for green, energy-efficient cooling, based on datacenter site climate features (free-cooling most of the year – more than 80% of time).

Unified system interconnect:

- Original interconnect (60 Gb/s) with 3-D torus topology with full MPI-2.0 support
- Global synchronization network

MPI 2.0 is supported for “SKIF Aurora” interconnect with optimized set of libraries. Some MPI collective operations offloaded to 3D torus NIC router and facilitated by synchronization and interrupt networks.

Hybrid computing — programmable accelerators in FPGA available in each “SKIF-Aurora” node, original approach featuring **AutoHDL** high-level FPGA programming. The **AutoHDL** enable expression of basic mathematical forms in structures, acceptable for FPGA design and fine-grain parallelism. Some **AutoHDL** implementations on “SKIF-Aurora” platform proved to outperform CPU computation speed by factor of 5x and more in floating-point operation intensive tasks.

Monitoring and management SKIF-Mon system harnesses multiple physical networks, including dedicated SKIF-ServNet sensors and management network. **SKIF-ServNet** has independent low-voltage power supply lines that provide access to system parameters even if the complete datacenter is powered down. **SKIF-Mon** system features graphical user interface that offers system administrator easy graphical analysis of entire supercomputer status and all of its components.



OTHER “SKIF-AURORA” COLLABORATORS

- Company Eurotech (Italy) — embedded and HPC computing solutions developer and provider (<http://www.eurotech.com>)
- Intel Corporation provides technical and marketing support to R&D alliance
- “AltLinux” provides a Linux distribution support with adaptation to needs of Aurora HPC platform
- “NICEVT” company participates with effort to support PGAS languages
- A number of Russian research organizations adapt and optimize applied software to “SKIF-Aurora” platform:
 - TESIS company: computational fluid dynamics code
 - Institute for Biochemical Physics RAS, Kintech company and others: open-source and proprietary codes in computational chemistry
 - Space Research Institute RAS: remote sensing data archive and transformation engine
 - Keldysh Applied Mathematics Institute RAS, Computational Mathematics and Cybernetics department of Lomonosov Moscow State University: electrostatics and computational physics codes
 - Other 30 SKIF-GRID collaborators

“SKIF-AURORA” INSTALLATION IN SOUTH URAL STATE UNIVERSITY

Active research community of South Ural State University demands constant increase in super-computing power available. “SKIF-Aurora” in SUSU installation — “SKIF-Aurora@SUSU”, — has been designed in collaborative effort by “SKIF-GRID” consortium.

- “SKIF-Aurora@SUSU” supercomputer setting up has been finished and evaluation period started, with peak performance of 24Tflops (<http://supercomputer.susu.ru>)
- “SKIF-Aurora@SUSU” holds 8th rank in Top50 rating of the most powerful computers of CIS countries (12 edition, March 2010)
- “SKIF-Aurora@SUSU” is the only Russian supercomputer installed outside Moscow, which is present in the TOP500. The “SKIF-Aurora@SUSU” holds the 450 ranks in this rating (34 edition, November 2009)

The “SKIF-Aurora@SUSU” has demonstrated high efficiency of real-life computational problems, which are typically addressed with help of SUSU supercomputer center resources. On some tasks the “SKIF-Aurora@SUSU” has demonstrated speedup of more than 3 times over previous generation cluster “SKIF Ural”.

SUSU SUPERCOMPUTER CENTER

SUSU has been created the most powerful and efficient Supercomputer Center in Ural and whole Eastern Russia regions which helps to solve fundamental and application tasks for the development of science potential and economy in that region — more than 100 projects have been done already not only for the local companies, but from other regions.

SUSU Supercomputer Center provides high-level services for an engineering research and analysis which are accessible through Internet to use in education framework and science research projects.

Key areas in which SUSU Supercomputer Center's services have been used and are most interested for:

- Mechanical engineering and manufacturing
- Steel industry
- Oil/Gas and Energy
- Textile industry
- Supercomputing and software development



“SKIF-AURORA” RUSSIAN PARTICIPANTS AND THEIR CONTRIBUTION



RSC SKIF — HPC SYSTEM INTEGRATOR AND SOLUTION PROVIDER FOR THE NEW LEVEL SUPERCOMPUTING
<http://www.rsc-skif.ru>

RSC SKIF — the only one company in Russia and CIS countries being developer, system integrator and provider of supercomputing solutions based on Intel architectures and technologies, state-of-the-art liquid cooling and broad scope of the own know-how. That company's potential helps:

- Deliver in practice the highest computing density based on industry standard x86 architecture processors
- Use a truly green design
- Provide highest reliability
- Get total silence in work of computing modules
- Provide full compatibility and guarantee of scalability based on Intel Cluster Ready certification
- Delivering in result the non-precedent low Total Cost of Ownership (TCO) and low power usage level

Professional team

RSC SKIF — the young and ambitious team of professionals with broad experience in technology R&D, project design, deployment of complex supercomputing solutions and integrated projects at enterprise wide scale as well as in practical usage and development of HPC applications.

Contribution in “SKIF-Aurora” project

RSC SKIF designed and implemented the whole integration project of “SKIF-Aurora” supercomputing system in South Ural State University including system design and deployment of the most efficient (green design) power supply and liquid cooling subsystems to provide PUE of the solution at industry-record low level — 1.2.



AILAMAZYAN PROGRAM SYSTEMS INSTITUTE OF THE RUSSIAN ACADEMY OF SCIENCES — AN IT-LEADER IN RUSSIA

<http://skif.pereslavl.ru/psi-info/index.en.html>



PSI RAS

Allamazyan Program Systems Institute of the Russian Academy of Sciences (PSI RAS) — a leading Russian research institution in the realm of information technologies,— was founded in 1984. Today's PSI RAS has a number of acknowledged achievements in the fields of high-end computer systems (supercomputers, grids), artificial intelligence, sensor networks, wide-area networks, distributed information systems and medical information systems. The PSI RAS was the head organization on the Russian side of the “SKIF-GRID” project. That means that in addition to research and development the institute was responsible for management and scientific supervision of the whole project.



Professional team

The PSI RAS employs 300 researchers and developers, among them one Academician (Full Member of RAS), one Corresponding Member of RAS, 18 Doctors of Sciences, 24 PhD, 240 specialists/MSc. In 1993 the PSI RAS initiated the establishment of Ailamazyan University of Pereslavl — the first Russian non-government university created in a small town. Today the university trains highly qualified specialists with professional knowledge and skills that are necessary for the PSI RAS. As result, now the major part of the PSI RAS staff is young talent and innovative people.

Contribution in “SKIF-Aurora” project

The following “SKIF-Aurora” subsystems and features (technical details are described above) were developed by PSI RAS:

- Original interconnect (60 Gb/s) with 3-D torus topology
 - Routing implemented in FPGA
 - MPI-2.0 with optimized collective operations
 - Effective support of non-MPI programming models for HPC — PGAS, GASNET, etc.
- Hybrid computing for floating-point operation intensive tasks — combining usage of standard i86-64 CPUs and FPGA-based accelerator
- Sensor and management network — SKIF-ServNet
- Monitoring and management system — SKIF-Mon

SOUTH URAL STATE UNIVERSITY



<http://www.susu.ac.ru/>

South Ural State University (SUSU, founded in 1943, Chelyabinsk city) — one of the best and biggest universities in Russian Federation, placing in Top10 Russia's Universities.

SUSU has 38 faculties, two special faculties for the additional education, Institute of the second education, 3 technical colleges, 13 branches in Russia and more than 55,000 students. SUSU provides education in all forms including distance learning technologies at scope of 300 highest education programs and 200 second and additional education programs.

In Professor and lecturer staff — more than 2,800 people, including 1,500 employees with PhD in different sciences. The total amount of research projects in SUSU is about 132.5 millions of rubles (more than \$4.4 millions).

South Ural State University has received the status of Russian National Research University on April 26, 2010 after participation in contest within 128 other Russia's universities.

10 years program of National Research University provides the possibility to create a world class level research university with dynamic development of sciences and technologies in the most priority directions as well as to get educated specialists for solving tasks in the energy and resource efficiency for high technology sectors of economy and in social area.

